ulm university universität ulm

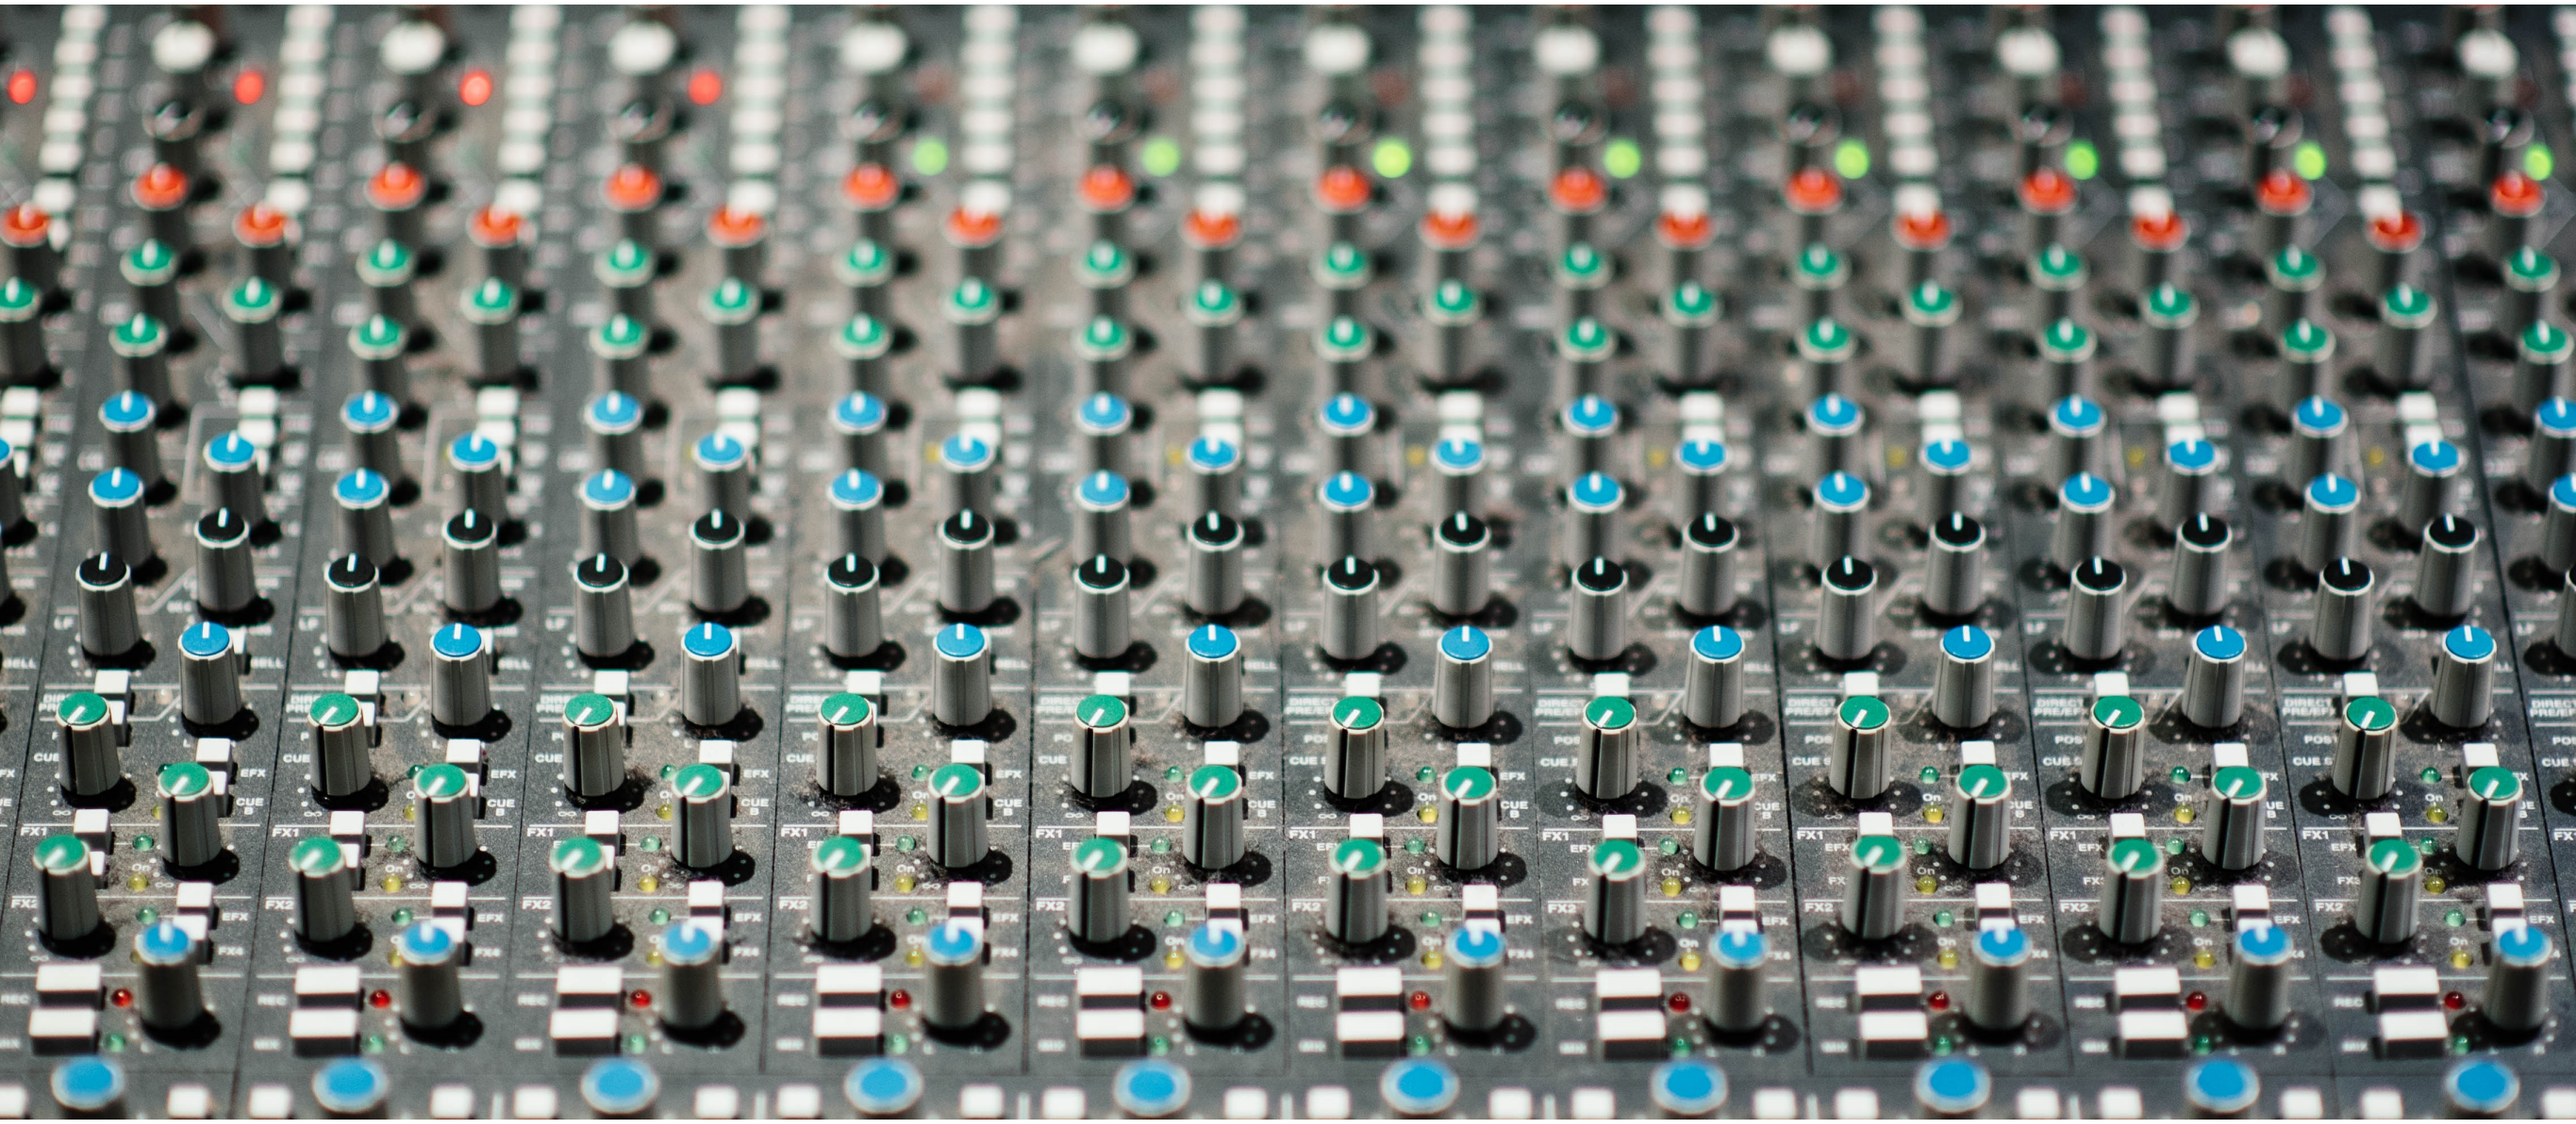

Scheduling Operating System Services

PhD Planner

Research Area: Operating Systems, Distributed Systems

Stefan Bonfert | Ulm University
Advisor: Stefan Wesner | Ulm University

What is this about?

- System calls are executed at arbitrary locations and thereby pollute caches and block shared data structures
- This limits performance of applications since they have to wait for system calls
- System calls can be executed remotely and asynchronously in, e.g., Microkernels
- CPU cores can be dedicated to run system calls. This reduces cache pollution and improves application speed

System Calls

- Interaction Between OS and Application
- Remote Execution in Microkernels
- Advantage:
  - Fewer Locking
  - Lower Cache Pollution

System Calls

- Interaction Between OS and Application
- Remote Execution in Microkernels
- Advantage:
  - Fewer Locking
  - Lower Cache Pollution

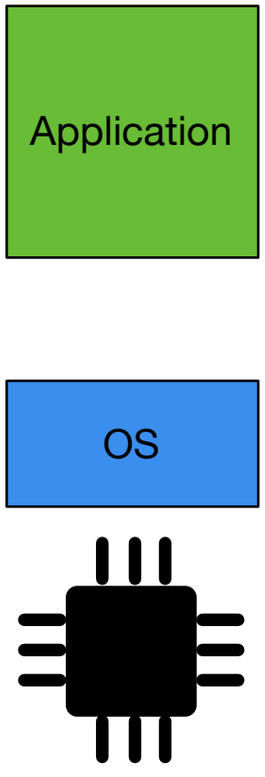

System Calls

- Interaction Between OS and Application
- Remote Execution in Microkernels
- Advantage:
  - Fewer Locking
  - Lower Cache Pollution

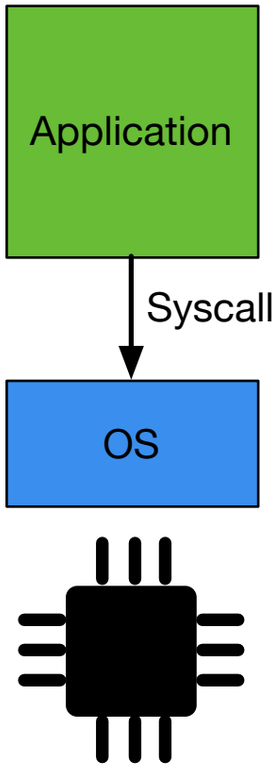

System Calls

- Interaction Between OS and Application
- Remote Execution in Microkernels
- Advantage:
  - Fewer Locking
  - Lower Cache Pollution

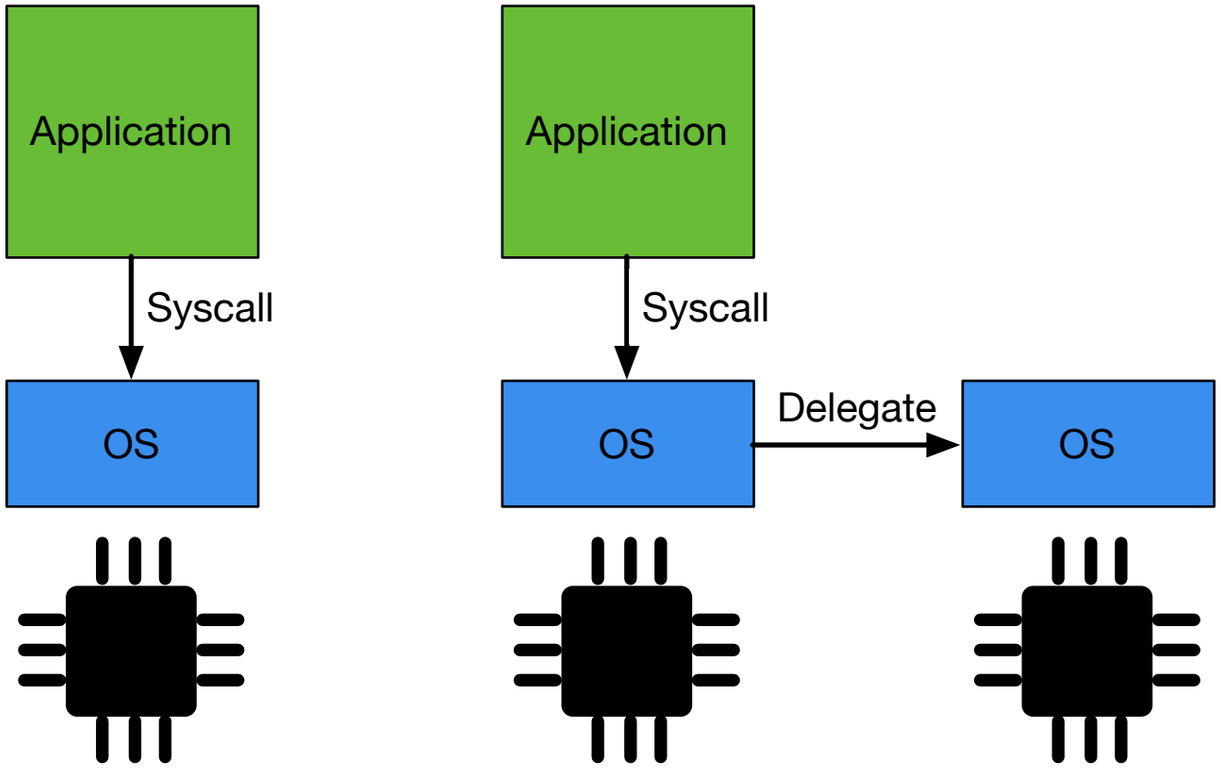

System Calls

- Interaction Between OS and Application
- Remote Execution in Microkernels
- Advantage:
  - Fewer Locking
  - Lower Cache Pollution

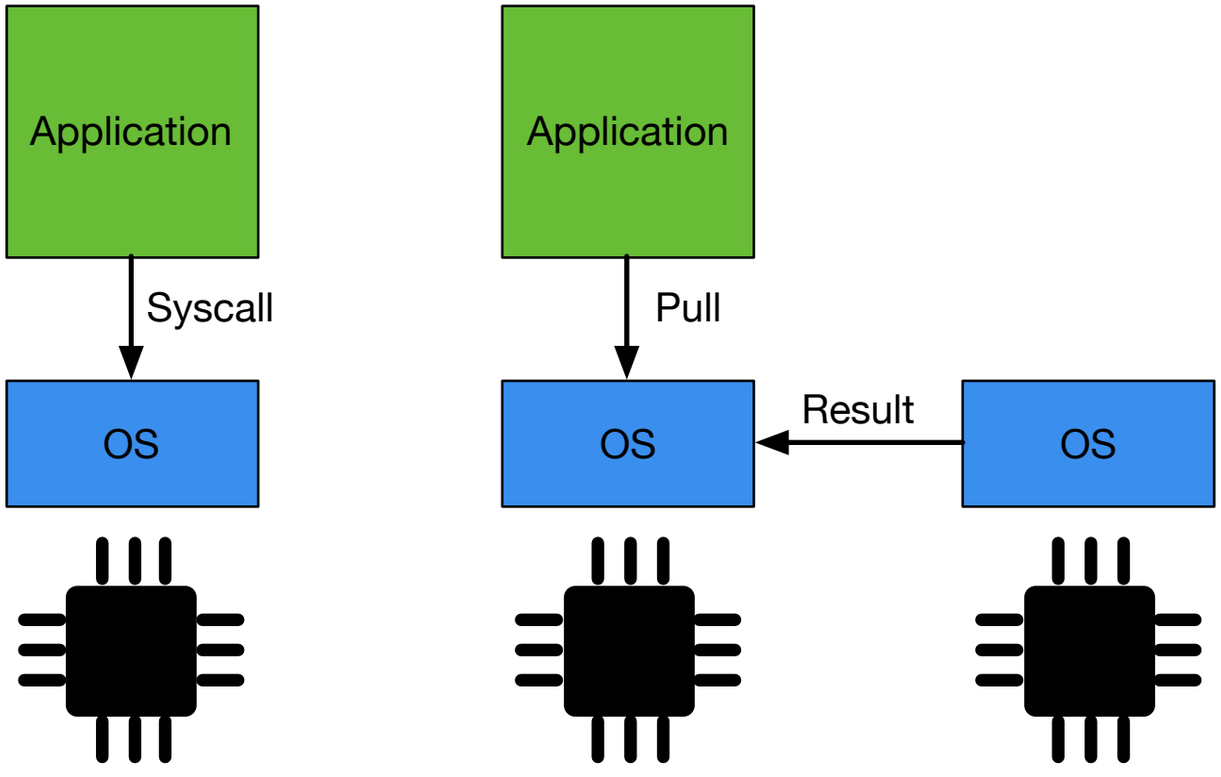

OS Services

- Specific System Call
- Offered by One or Multiple Cores
- Replication:
  - Improves Locality
  - Synchronisation & Consistency

OS Services

- Specific System Call
- Offered by One or Multiple Cores
- Replication:
  - Improves Locality
  - Synchronisation & Consistency

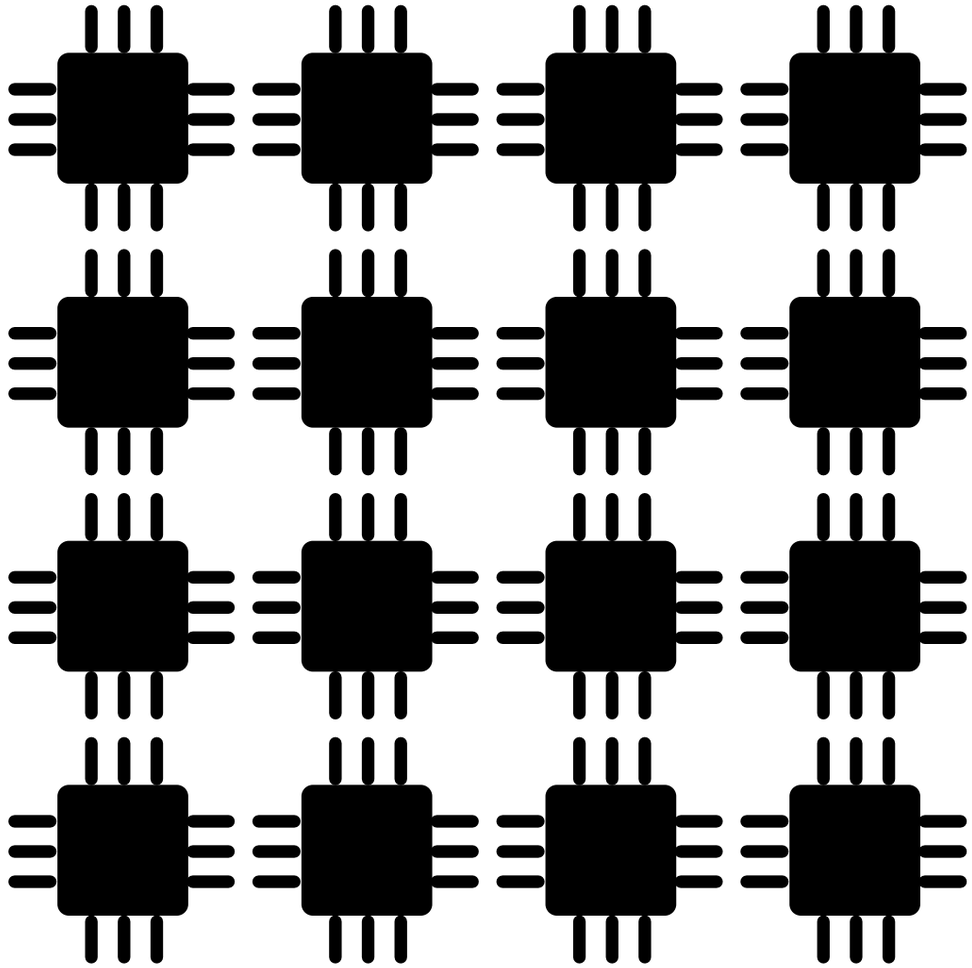

OS Services

- Specific System Call
- Offered by One or Multiple Cores
- Replication:
  - Improves Locality
  - Synchronisation & Consistency

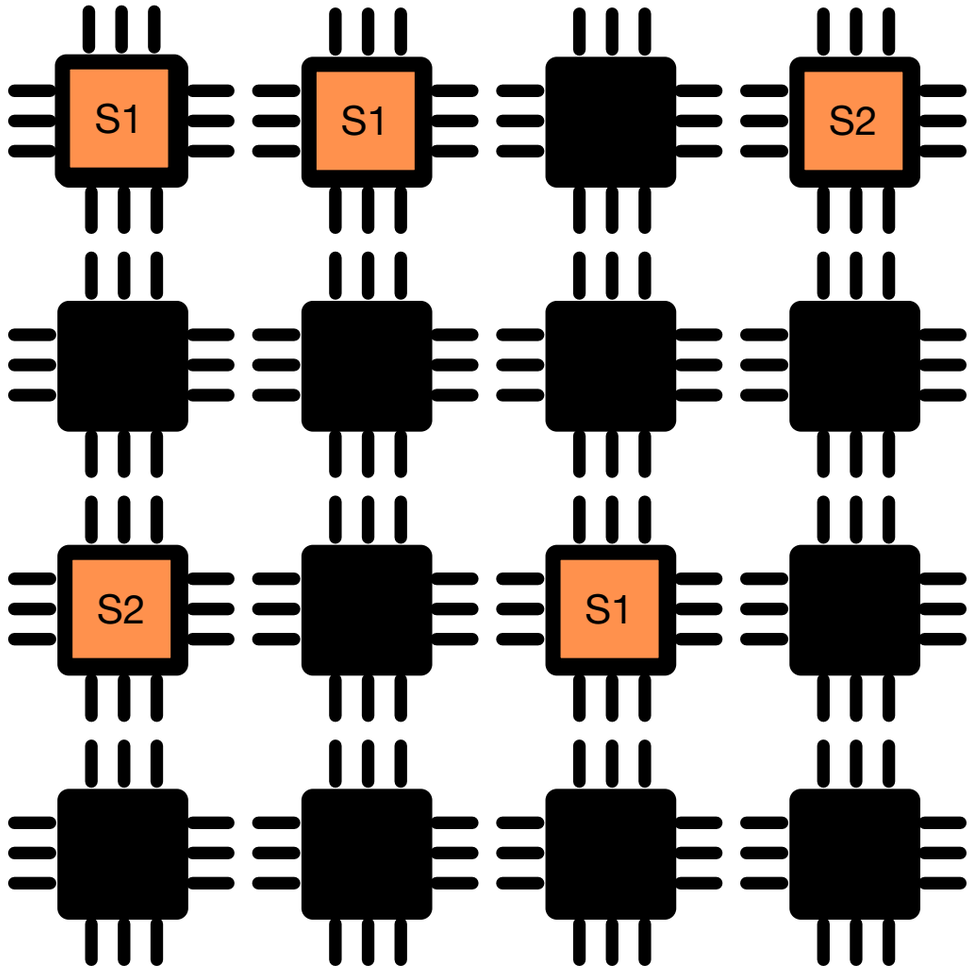

OS Services

- Specific System Call
- Offered by One or Multiple Cores
- Replication:
  - Improves Locality
  - Synchronisation & Consistency

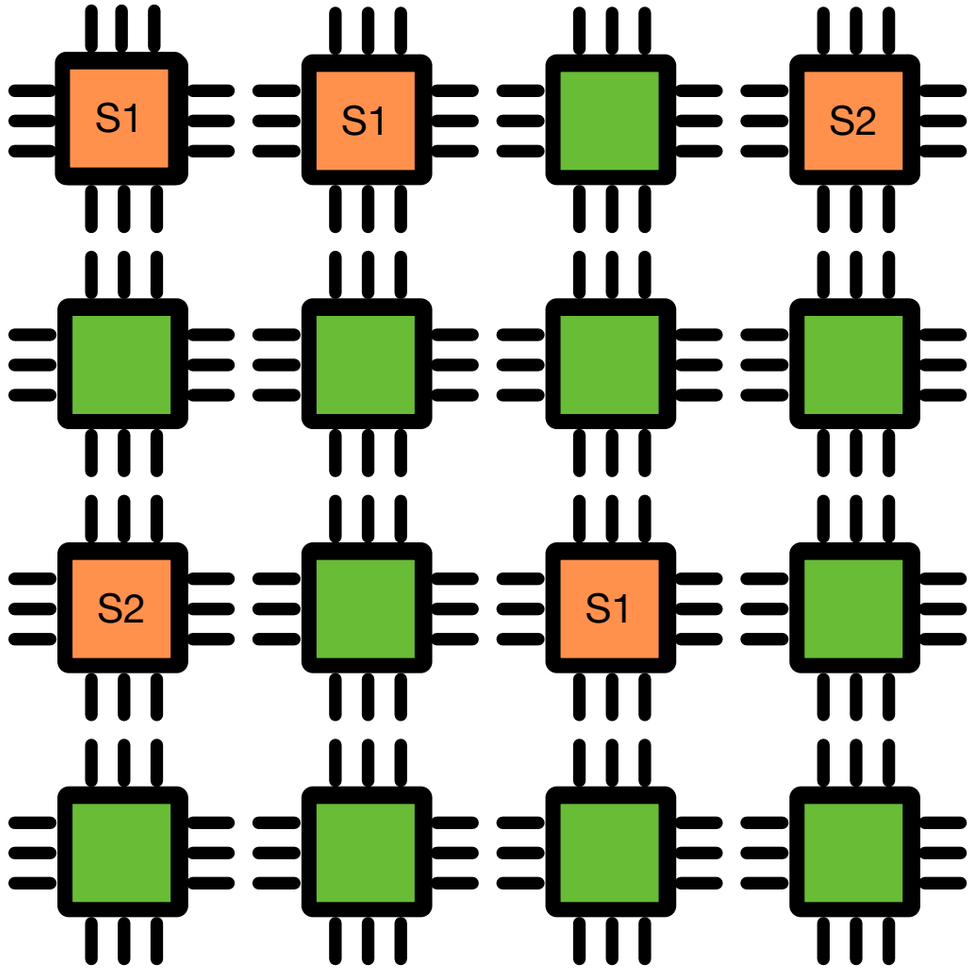

Scope

- Multi-Node, Multi-Application System
- HPC Applications
  - High Concurrency
  - Recurring Tasks
  - No User Interaction
  - Non-Preemptible Execution

Approach

- Monitor System Calls
- Extract Affinities of Tasks to Services
- Cost Model
  - Communication to Services
  - Replication of Services
    - Synchronisation & Consistency
  - Location of Input Data
  - Resource Utilisation
- Hierarchical Scheduler for Tasks
- Co-Schedule System Services

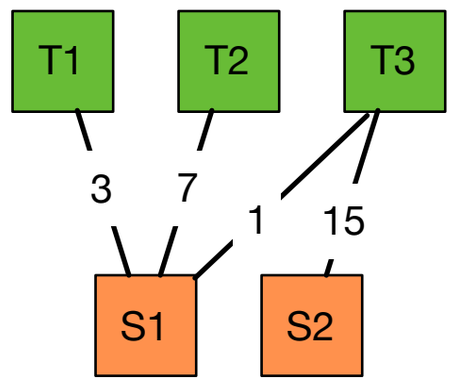

Research Questions

- How does OS service placement affect application performance?
- How do scheduling of application tasks and OS services influence each other?
- What is the optimal placement strategy for OS services to achieve maximum application performance?
- What is the reconfiguration cost for different OS services?

Questions?